Gaussian Process Quadrature Moment Transform

Jakub Prüher, Ondřej Straka

January 6, 2017

arXiv:1701.01356v1 [stat.ME] 5 Jan 2017

Abstract

Computation of moments of transformed random variables is a problem appearing in many engineering applications. The current methods for moment transformation are mostly based on the classical quadrature rules which cannot account for the approximation errors. Our aim is to design a method for moment transformation for Gaussian random variables which accounts for the error in the numerically computed mean. We employ an instance of Bayesian quadrature, called Gaussian process quadrature (GPQ), which allows us to treat the integral itself as a random variable, where the integral variance informs about the incurred integration error. Experiments on the coordinate transformation and nonlinear filtering examples show that the proposed GPQ moment transform performs better than the classical transforms.

1 Introduction

The goal of a moment transformation is to compute statistical moments of random variables transformed through an arbitrary nonlinear function. The moment transforms find their uses in such problems as sensor system design [Zangl and Steiner, 2008], optimal control [Heine et al., 2006; Ross et al., 2015] and they are also an indispensable part of local nonlinear filters and smoothers [Duník et al., 2013; Sandblom and Svensson, 2012; Särkkä et al., 2016; Tronarp et al., 2016; Wu et al., 2006]. These algorithms estimate a state of dynamical systems based on noisy measurements and are applied in solving a broad array of engineering problems such as aircraft guidance [Smith et al., 1962], GPS navigation [Grewal et al., 2007], weather forecasting [Gillijns et al., 2006], telecommunications [Jiang et al., 2003] and finance [Bhar, 2010] to name a few.

Recursive nonlinear filters can be divided into two categories: global and local. Particle filters are typical representatives of the global filters, which are characterized by weaker assumptions and a higher computational demand. On the contrary, the local filters trade off more limiting assumptions for computational simplicity. For tractability reasons, the local filters often leverage the joint Gaussianity assumption of state and measurement. The main problem then lies in computation of transformed means and covariances, which are subsequently combined with a measurement in an update rule to produce the filtering state estimate. Examples of well-known local nonlinear filters include the unscented Kalman filter (UKF) [Julier et al., 2000], the cubature Kalman filter (CKF) [Arasaratnam and Haykin, 2009] and the Gauss-Hermite Kalman filter (GHKF) [Wu et al., 2006], which are collectively known as *sigma-point* filters, and which are characterized by their reliance on the classical numerical integration schemes.

A limitation of classical integral approximations, such as the Gauss-Hermite quadrature, is that they are specifically designed to achieve zero error on a narrow class of functions (typically polynomials up to a given degree). Since many nonlinearities encountered in practical problems do not fall into this category, the integrals are, more often than not, approximated with errors which go unaccounted for. Even though error estimates for the classical rules exist, their computation is tedious in practice as they require higher order derivatives [Gautschi, 2004].

In recent years, the Bayesian quadrature (BQ) has been gaining attention as an exciting alternative for approximate evaluation of integrals. According to Diaconis [1988], the origins of this method lie as far as Poincaré's publication [Poincaré, 1896] from 1896. Later developments came with O'Hagan's work on Bayes-Hermite quadrature O'Hagan [1991] and the Bayesian Monte Carlo [Rasmussen and Ghahramani, 2003]. Seeing

BQ as an instance of a probabilistic approach to numerical computing spawned an emerging field of probabilistic numerics with its share of contributions [Briol et al., 2015; Oates et al., 2015; Osborne and Garnett, 2012]. Contrary to the classical rules, the BQ minimizes an average error on a wider class of functions [Briol et al., 2015; Minka, 2000]. The numerical integration process is treated as a problem of Bayesian inference, where, in accordance with the Bayesian paradigm, an integral prior is transformed into a posterior by conditioning on obtained evaluations of the integrated function. The result of integration is much more informative, because it is no longer a single value, but an entire distribution. The posterior mean estimates the integral value, whereas the posterior variance can be construed as a model of the integration error. The BQ approach is uniquely suited to the purpose of our article, which is accounting for the integration errors in the moment transform process.

Application of the BQ in nonlinear sigma-point filtering was first investigated by Särkkä et al. [2014], where the authors elucidate connections between the BQ and the classical rules, but their algorithms do not make use of the integral variance. When our GPQ moment transform is applied in the nonlinear sigma-point filtering, it leads to a similar approach to the previously proposed Gaussian process assumed density filter (GP-ADF) [Deisenroth et al., 2009] and the Gaussian process unscented Kalman filter (GP-UKF) [Ko et al., 2007]. Both GP-ADF and GP-UKF use the GP models for system identification that takes place prior to running the filters. In our case, however, the crucial difference lies in the fact that the resulting GPQ filters do not require a system identification phase.

In our previous contribution [Prüher and Šimandl, 2016], we showed how to incorporate integral variance into a nonlinear sigma-point filtering algorithm. This article crystallizes the results of the previous publication into a widely applicable general GPQ moment transform. Applications in sigma-point filtering are enriched with a target tracking example and an additional numerical experiments on common nonlinear coordinate transformation are provided. We also further analyse properties of the proposed transform and give a theoretical proof for positive semi-definiteness of the resulting covariance matrix.

The rest of this article is organized as follows. Section 2 introduces the general problem of the classical quadrature based moment transforms. Section 3 outlines the Gaussian process quadrature (GPQ). In Section 4, we describe the proposed moment transform based on the GPQ. Numerical experiments and performance evaluations are in Section 5. Finally, Section 6 concludes the article.

2 Problem Statement

Consider an input Gaussian random variable $\mathbf{x} \in \mathbb{R}^D$ with mean $\mathbf{m}$ and covariance $\mathbf{P}$ transformed through a nonlinear vector function

$$\mathbf{y} = \mathbf{g}(\mathbf{x}) \qquad \mathbf{x} \sim \mathrm{N}(\mathbf{m}, \mathbf{P}) \tag{2.1}$$

producing an output $\mathbf{y} \in \mathbb{R}^E$. Even though the output is not Gaussian (due to the nonlinearity of $\mathbf{g}$), we approximate both variables as jointly Gaussian distributed, so that

$$\begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} \sim \mathrm{N}\left(\begin{bmatrix} \mathbf{m} \\ \boldsymbol{\mu} \end{bmatrix}, \begin{bmatrix} \mathbf{P} & \mathbf{C} \\ \mathbf{C}^\top & \boldsymbol{\Pi} \end{bmatrix} \right), \tag{2.2}$$

where the transformed moments are given by the following Gaussian weighted integrals

$$\boldsymbol{\mu} = \mathbb{E}_{\mathbf{x}}[\mathbf{g}(\mathbf{x})] = \int \mathbf{g}(\mathbf{x}) \mathrm{N}(\mathbf{x} \,|\, \mathbf{m}, \mathbf{P}) \,\mathrm{d}\mathbf{x}, \tag{2.3}$$

$$\boldsymbol{\Pi} = \mathbb{C}_{\mathbf{x}}[\mathbf{g}(\mathbf{x})] = \int (\mathbf{g}(\mathbf{x}) - \boldsymbol{\mu})(\mathbf{g}(\mathbf{x}) - \boldsymbol{\mu})^\top \mathrm{N}(\mathbf{x} \,|\, \mathbf{m}, \mathbf{P}) \,\mathrm{d}\mathbf{x}, \tag{2.4}$$

$$\mathbf{C} = \mathbb{C}_{\mathbf{x}}[\mathbf{x}, \mathbf{g}(\mathbf{x})] = \int (\mathbf{x} - \mathbf{m})(\mathbf{g}(\mathbf{x}) - \boldsymbol{\mu})^\top \mathrm{N}(\mathbf{x} \,|\, \mathbf{m}, \mathbf{P}) \,\mathrm{d}\mathbf{x}, \tag{2.5}$$

where $\mathrm{N}(\mathbf{x}\,|\,\mathbf{m},\,\mathbf{P})$ denotes probability density of Gaussian random variable. The goal of a moment transform is to compute the moments in eqs. (2.3) to (2.5) given the moments of the input variable. Since $\mathbf{g}$ is nonlinear, the integrals cannot be solved analytically in general and have to be approximated by a numerical quadrature. An integral w.r.t. a Gaussian with arbitrary mean and covariance can be converted to an integral over a standard Gaussian, such that

$$\int \mathbf{g}(\mathbf{x})\mathrm{N}(\mathbf{x}\,|\,\mathbf{m},\,\mathbf{P})\,\mathrm{d}\mathbf{x} = \int \mathbf{g}(\mathbf{m}+\mathbf{L}\boldsymbol{\xi})\mathrm{N}(\boldsymbol{\xi}\,|\,\mathbf{0},\,\mathbf{I})\,\mathrm{d}\boldsymbol{\xi}, \tag{2.6}$$

where we used a change of variables $\mathbf{x} = \mathbf{m} + \mathbf{L}\boldsymbol{\xi}$ with $\mathbf{P} = \mathbf{L}\mathbf{L}^\top$. The standard numerical integration rules can now be applied, which leads to a weighted sum approximation

$$\int \mathbf{g}(\mathbf{m}+\mathbf{L}\boldsymbol{\xi})\mathrm{N}(\boldsymbol{\xi}\,|\,\mathbf{0},\,\mathbf{I})\,\mathrm{d}\boldsymbol{\xi} = \sum_{i=1}^{N} w_i \mathbf{g}(\mathbf{m}+\mathbf{L}\boldsymbol{\xi}_i), \tag{2.7}$$

where w_i are the quadrature weights and $\mathbf{x}_i = \mathbf{m} + \mathbf{L}\boldsymbol{\xi}_i$ are the sigma-points (evaluation points, design points, abscissas), both of which are prescribed by the specific quadrature rule to satisfy certain optimality criteria. The vectors $\boldsymbol{\xi}_i$ are called unit sigma-points.

For instance, the r-th order Gauss-Hermite (GH) rule [Gautschi, 2004; Ito and Xiong, 2000] uses sigma-points, which are determined as the roots of the r-th degree univariate Hermite polynomial $H_r(x)$. Integration of vector valued functions ($D > 1$) is handled by a multidimensional grid of points formed by the Cartesian product, leading to the exponential growth ($N = r^D$) w.r.t. dimension D. The GH weights are computed according to [Särkkä, 2013] as

$$w_i = \frac{r!}{[rH_{r-1}(x^{(i)})]^2} \ . \tag{2.8}$$

The rule incurs no integration error if the integrand is a (multivariate) polynomial of *pseudo-degree* $\leq 2r-1$. The well-known Unscented transform (UT) [Julier et al., 2000] is also a simple quadrature rule, that uses $N = 2D+1$ deterministically chosen sigma-points $\mathbf{x}_i = \mathbf{m} + \mathbf{L}\boldsymbol{\xi}_i$ with unit sigma-points defined as columns of the matrix

$$\begin{bmatrix}\boldsymbol{\xi}_0 & \boldsymbol{\xi}_1 & \dots & \boldsymbol{\xi}_{2D}\end{bmatrix} = \begin{bmatrix}\mathbf{0} & c\mathbf{I}_D & -c\mathbf{I}_D\end{bmatrix} \tag{2.9}$$

where $\mathbf{I}_D$ denotes $D \times D$ identity matrix. The corresponding UT weights are defined by

$$w_0 = \frac{\kappa}{D+\kappa}, \quad w_i = \frac{1}{2(D+\kappa)}, \quad i = 1,\dots,2D \tag{2.10}$$

with scaling factor $c = \sqrt{D+\kappa}$. The UT rule can integrate (multivariate) polynomials of *total degree* ≤ 3 without incurring approximation error. Spherical-radial (SR) integration rule, which is a basis of the CKF [Arasaratnam and Haykin, 2009], is very similar to the UT, but lacks the centre point. Thus it uses $N = 2D$ sigma-points given by

$$\begin{bmatrix}\boldsymbol{\xi}_1 & \dots & \boldsymbol{\xi}_{2D}\end{bmatrix} = \begin{bmatrix}c\mathbf{I}_D & -c\mathbf{I}_D\end{bmatrix} \tag{2.11}$$

with $c = \sqrt{D}$ and weights $w_i = 1/2D, \quad i = 1,\dots,2D$. The UT and SR rule are all instances of the fully symmetric rules [McNamee and Stenger, 1967]. Together with GH all of these are examples of classical numerical quadratures.

In the next section, we introduce an alternative view of numerical integration, upon which, we base our proposed moment transform.

3 Gaussian Process Quadrature

The main difference between the classical quadrature and the Bayesian quadrature, is that in the Bayesian case the whole numerical procedure is viewed as a probabilistic inference, where, after conditioning on the data, an

entire density over the solution is obtained as a result (rather than a single value). That is to say, we start with a prior distribution over the integral which is then transformed into a posterior distribution given our sigma-points and function evaluations. Prior over the integral is induced by putting a prior on the integrated function itself.

3.1 Gaussian Process Regression Model

Uncertainty over functions is naturally expressed by a stochastic process. In Bayesian quadrature, Gaussian processes (GP) are used for their favourable analytical properties. Gaussian process is a collection of random variables indexed by elements of an index set, any finite number of which has a joint Gaussian density [Rasmussen and Williams, 2006]. That is, for any finite set of indices $\mathbf{X}' = \begin{bmatrix}\mathbf{x}'_1 & \dots & \mathbf{x}'_m\end{bmatrix}$, it holds that

$$\begin{bmatrix}g(\mathbf{x}'_1) & \dots & g(\mathbf{x}'_m)\end{bmatrix}^\top \sim \mathrm{N}(\mathbf{0}, \mathbf{K}) \tag{3.1}$$

where the kernel (covariance) matrix $\mathbf{K}$ is made up of pair-wise evaluations of the kernel (covariance) function. The element of $\mathbf{K}$ at position (i, j) is given by $[\mathbf{K}]_{ij} = k(\mathbf{x}_i, \mathbf{x}_j) = \mathbb{C}_g[g(\mathbf{x}_i), g(\mathbf{x}_j)]$, where $\mathbb{C}$ is the covariance operator. Choosing a kernel, which in principle can be any symmetric positive definite function of two arguments, introduces modelling assumptions about the underlying function. Since the GP regression is a non-parametric model, it is more expressive than a parametric fixed-order polynomial regression models used in construction of the classical quadrature rules. In Bayesian terms, choosing a kernel specifies a GP prior $p(g)$ over functions. Updating the prior with the data $\mathcal{D} = \{(\mathbf{x}_i,\ g(\mathbf{x}_i))\}_{i=1}^N$, which consist of the sigma-points $\mathbf{X} = \begin{bmatrix}\mathbf{x}_1 & \dots & \mathbf{x}_N\end{bmatrix}$ and the function evaluations $\mathbf{y} = \begin{bmatrix}g(\mathbf{x}_1) & \dots & g(\mathbf{x}_N)\end{bmatrix}^\top$, leads to a GP posterior $p(g \mid \mathcal{D})$ with predictive mean and variance given by

$$\mathbb{E}_g[g(\mathbf{x}) \mid \mathcal{D}] = m_g(\mathbf{x}) = \mathbf{k}^\top(\mathbf{x})\mathbf{K}^{-1}\mathbf{y}, \tag{3.2}$$

$$\mathbb{V}_g[g(\mathbf{x}) \mid \mathcal{D}] = \sigma_g^2(\mathbf{x}) = k(\mathbf{x}, \mathbf{x}) - \mathbf{k}^\top(\mathbf{x})\mathbf{K}^{-1}\mathbf{k}(\mathbf{x}), \tag{3.3}$$

where the i-th element of $\mathbf{k}(\mathbf{x})$ is $[\mathbf{k}(\mathbf{x})]_i = k(\mathbf{x}, \mathbf{x}_i)$ and $\mathbb{V}$ denotes the variance operator. These simple equations follow from the formula for conditional Gaussian densities [Rasmussen and Williams, 2006]. Figure 1 depicts predictive moments of the GP posterior. Notice, that in places where the function evaluations are lacking, the GP model is more uncertain.

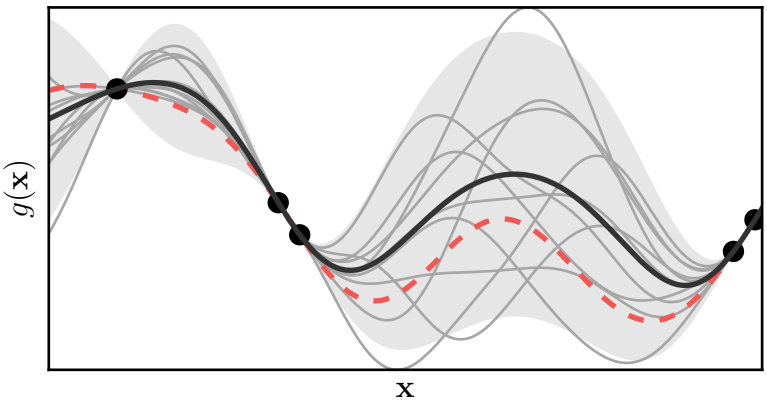

Figure 1: True function (dashed), GP posterior mean (solid), observed function values (dots) and GP posterior samples (grey). The shaded area represents GP posterior predictive uncertainty ($\pm 2\,\sigma_g(\mathbf{x})$), which is collapsed near the observations.

At first, introducing randomness[1] over a *known*[2] integrand may seem quite counter-intuitive. However, consider the fact that the quadrature rule of the form (2.7) only sees the integrand through a limited number

[1] Here we mean *epistemic uncertainty*, which is due to the lack of knowledge not due to inherent randomness.
[2] In a sense, that it can be evaluated for any argument.

of function values, which effectively means, that the rule is unaware of the function behaviour in areas where no evaluations are available. The GP regression model then serves as an instrument allowing us to acknowledge this reality.

3.2 Integral Moments

Using a GP for modelling the integrand has a great analytical advantages. Note that, since we use a GP, the posterior density over the integral is Gaussian, which is due to the fact that an integral is a linear operator acting on a Gaussian distributed function.[3] The posterior mean and variance of the integral $\mathbb{E}_{\mathbf{x}}[g(\mathbf{x})] = \int g(\mathbf{x})p(\mathbf{x})\,\mathrm{d}\mathbf{x}$ are [Rasmussen and Ghahramani, 2003]

$$\mathbb{E}_g[\mathbb{E}_{\mathbf{x}}[g(\mathbf{x})] \mid \mathcal{D}] = \mathbb{E}_{\mathbf{x}}[\mathbb{E}_g[g(\mathbf{x}) \mid \mathcal{D}]] = \mathbb{E}_{\mathbf{x}}\big[\mathbf{k}^\top(\mathbf{x})\big]\mathbf{K}^{-1}\mathbf{y}, \tag{3.4}$$

$$\begin{aligned}\mathbb{V}_g[\mathbb{E}_{\mathbf{x}}[g(\mathbf{x})] \mid \mathcal{D}] &= \mathbb{E}_{\mathbf{x},\mathbf{x}'}\big[\mathbb{C}_g\big[g(\mathbf{x}), g(\mathbf{x}') \mid \mathcal{D}\big]\big] \\ &= \mathbb{E}_{\mathbf{x},\mathbf{x}'}\big[k(\mathbf{x},\mathbf{x}')\big] - \mathbb{E}_{\mathbf{x}}\big[\mathbf{k}^\top(\mathbf{x})\big]\mathbf{K}^{-1}\mathbb{E}_{\mathbf{x}'}\big[\mathbf{k}(\mathbf{x}')\big].\end{aligned} \tag{3.5}$$

From eq. (3.4), we can see that the mean of the integral is identical to integrating the GP posterior mean, which effectively serves as an approximation of the integrated function. The posterior variance of the integral, given by eq. (3.5), can be construed as a model of the integration error. The Figure 2 depicts the schematic view of the GP quadrature, where integral is applied to a GP distributed integrand yielding a Gaussian density over the value of the integral. Alternatively, one could imagine integrating each realization of the GP posterior separately, yielding a different result every time.

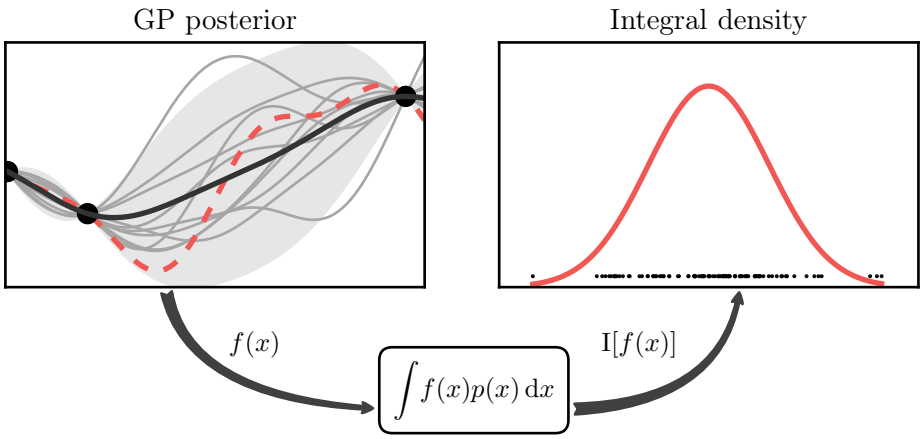

Figure 2: Gaussian process quadrature. GP distributed function is mapped through a linear operator to yield a Gaussian density over its solution. The GP posterior mean approximation (solid black) of the true function (dashed) and the GP predictive variance (gray band) are based on the function evaluations (dots).

In light of this view, we could think of the classical quadratures as returning a Dirac distribution over the solution, where all the probability mass is concentrated at one value. In this regard, Bayesian quadrature rule is more realistic, because it is able to acknowledge uncertainty in the integrand due to the minimal number of available evaluations.

Also worth noting, is that the classical quadrature rules define precise locations of sigma-points, whereas the BQ does not prescribe any point sets, which raises questions about their placement. In [Minka, 2000; O'Hagan, 1991], the optimal point set is determined by minimizing the posterior integral variance (3.5). Another approach developed in [Osborne et al., 2012] uses a sequential active sampling scheme. Applications of GPQ in this article rely on point sets of the classical rules.

[3] This is a generalization of the invariance property of Gaussians under affine transformations.

4 GPQ Moment Transformation

In this section, we propose the moment transform based on the GPQ and show how it implicitly utilizes posterior integral variance in the moment transformation process. First, we define a general GPQ moment transform, which is applicable for any kernel function, and then give relations for a concrete transform based on the popular RBF (Gaussian) kernel.

We begin by employing a GPQ for approximate evaluation of the moment integrals in eqs. (2.3) to (2.5). For a moment, consider a case when the nonlinearity in eq. (2.1) is a scalar function $g(\mathbf{x}) : \mathbb{R}^D \to \mathbb{R}$. Since the source of uncertainty is now, not only in the input $\mathbf{x}$, but the nonlinearity g as well, the transformed moments also need to reflect this fact. The GPQ transform then approximates the moments as follows

$$\mathbb{E}[y] = \mathbb{E}_{\mathbf{x}}[g(\mathbf{x})] \approx \mathbb{E}_{g,\mathbf{x}}[g(\mathbf{x})] \tag{4.1}$$
$$\mathbb{V}[y] = \mathbb{V}_{\mathbf{x}}[g(\mathbf{x})] \approx \mathbb{V}_{g,\mathbf{x}}[g(\mathbf{x})] \tag{4.2}$$
$$\mathbb{C}[\mathbf{x}, y] = \mathbb{C}_{\mathbf{x}}[\mathbf{x}, g(\mathbf{x})] \approx \mathbb{C}_{g,\mathbf{x}}[\mathbf{x}, g(\mathbf{x})] \tag{4.3}$$

where, using the law of total expectation and variance, we can further write

$$\mathbb{E}_{g,\mathbf{x}}[g(\mathbf{x})] = \mathbb{E}_g[\mathbb{E}_{\mathbf{x}}[g(\mathbf{x})]] = \mathbb{E}_{\mathbf{x}}[\mathbb{E}_g[g(\mathbf{x})]], \tag{4.4}$$
$$\mathbb{V}_{g,\mathbf{x}}[g(\mathbf{x})] = \mathbb{E}_g[\mathbb{V}_{\mathbf{x}}[g(\mathbf{x})]] + \mathbb{V}_g[\mathbb{E}_{\mathbf{x}}[g(\mathbf{x})]] \tag{4.5}$$
$$= \mathbb{E}_{\mathbf{x}}[\mathbb{V}_g[g(\mathbf{x})]] + \mathbb{V}_{\mathbf{x}}[\mathbb{E}_g[g(\mathbf{x})]], \tag{4.6}$$
$$\mathbb{C}_{g,\mathbf{x}}[\mathbf{x}, g(\mathbf{x})] = \mathbb{E}_{\mathbf{x}}[\mathbf{x}\mathbb{E}_g[g(\mathbf{x})]] - \mathbb{E}_{\mathbf{x}}[\mathbf{x}]\mathbb{E}_{g,\mathbf{x}}[g(\mathbf{x})]. \tag{4.7}$$

The eq. (4.4) shows that the mean of the integral is equivalent to integrating the GP mean function. Since the variance decompositions in eqs. (4.5) to (4.6) are equivalent, both can be used to achieve the same goal. The form (4.6) was utilized in derivation of the GP-ADF [Deisenroth et al., 2012], which relies on the solution to the problem of prediction with GPs at uncertain inputs [Girard et al., 2003]. So, even though these results were derived to solve a seemingly different problem, we point out, that by using the form (4.6), the uncertainty of the mean integral (as seen in the last term of eq. (4.5)) is implicitly reflected in the resulting covariance. Furthermore, the form (4.6) is preferable, because it is more amenable to analytical expression and implementation. Note, that for the deterministic case, when the *integrand* variance $\mathbb{V}_g[g(\mathbf{x})] = 0$ and the *integral* variance $\mathbb{V}_g[\mathbb{E}_{\mathbf{x}}[g(\mathbf{x})]] = 0$, the eqs. (4.4) to (4.7) fall back to the classical expressions given by eqs. (2.3) to (2.5). Compared to the deterministic case the transformed GPQ variance is inflated by the uncertainty in g.

4.1 Derivations of the transformed moments

So far we have stated the results only for the case of scalar function. In the following summary, general vector functions $\mathbf{g}(\mathbf{x}) : \mathbb{R}^D \to \mathbb{R}^E$ are modelled by a single GP. That is, one GP models every output dimension using the same values of kernel parameters. Note, that in the following derivations we omit the conditioning on data in the GP predictive moments and use the shorthand $\mathbb{E}_{\mathbf{g}}[\mathbf{g}(\mathbf{x})] \triangleq \mathbb{E}_{\mathbf{g}}[\mathbf{g}(\mathbf{x}) \mid \mathcal{D}]$, $\mathbb{C}_{\mathbf{g}}[\mathbf{g}(\mathbf{x})] \triangleq \mathbb{C}_{\mathbf{g}}[\mathbf{g}(\mathbf{x}) \mid \mathcal{D}]$.

Expressions for GPQ transformed moments are derived by plugging in the GP predictive moments from eqs. (3.2) to (3.3) into the general expressions in eqs. (4.4) to (4.7). The transformed mean in eq. (4.4) thus becomes

$$\boldsymbol{\mu}_{\mathrm{G}} = \mathbb{E}_{\mathbf{x}}[\mathbb{E}_{\mathbf{g}}[\mathbf{g}(\mathbf{x})]] = \mathbf{Y}^\top \mathbf{K}^{-1} \mathbb{E}_{\mathbf{x}}\big[\mathbf{k}(\mathbf{x})\big] = \mathbf{Y}^\top \mathbf{w}. \tag{4.8}$$

The transformed covariance in eq. (4.6) can be decomposed

$$\boldsymbol{\Pi}_{\mathrm{G}} = \mathbb{C}_{\mathbf{g},\mathbf{x}}[\mathbf{g}(\mathbf{x})] = \mathbb{C}_{\mathbf{x}}[\mathbb{E}_{\mathbf{g}}[\mathbf{g}(\mathbf{x})]] + \mathbb{E}_{\mathbf{x}}[\mathbb{C}_{\mathbf{g}}[\mathbf{g}(\mathbf{x})]] \tag{4.9}$$
$$= \mathbb{E}_{\mathbf{x}}\big[\mathbb{E}_{\mathbf{g}}[\mathbf{g}(\mathbf{x})]\mathbb{E}_{\mathbf{g}}[\mathbf{g}(\mathbf{x})]^\top\big] - \boldsymbol{\mu}_{\mathrm{G}}\boldsymbol{\mu}_{\mathrm{G}}^\top + \mathbb{E}_{\mathbf{x}}[\mathbb{C}_{\mathbf{g}}[\mathbf{g}(\mathbf{x})]] \tag{4.10}$$
$$= \mathbf{Y}^\top \mathbf{K}^{-1} \mathbb{E}_{\mathbf{x}}\big[\mathbf{k}(\mathbf{x})\mathbf{k}^\top(\mathbf{x})\big] \mathbf{K}^{-1}\mathbf{Y} - \boldsymbol{\mu}_{\mathrm{G}}\boldsymbol{\mu}_{\mathrm{G}}^\top + \sigma^2\mathbf{I}, \tag{4.11}$$

where $\sigma^2 = \mathbb{E}_{\mathbf{x}}\big[\sigma_g^2(\mathbf{x})\big] = \mathbb{E}_{\mathbf{x}}[k(\mathbf{x},\mathbf{x})] - \mathrm{tr}\left(\mathbb{E}_{\mathbf{x}}\big[\mathbf{k}(\mathbf{x})\mathbf{k}^\top(\mathbf{x})\big]\mathbf{K}^{-1}\right)$. The diagonal matrix in the last term reflects the fact that the outputs of $\mathbf{g}$ are not correlated (modelled independently). Finally, the cross-covariance becomes

$$\mathbf{C}_{\mathrm{G}} = \mathbb{C}_{\mathbf{g},\mathbf{x}}[\mathbf{x},\mathbf{g}(\mathbf{x})] = \mathbb{E}_{\mathbf{x}}\big[\mathbf{x}\mathbb{E}_{\mathbf{g}}[\mathbf{g}(\mathbf{x})]^\top\big] - \mathbb{E}_{\mathbf{x}}[\mathbf{x}]\mathbb{E}_{\mathbf{x}}[\mathbb{E}_{\mathbf{g}}[\mathbf{g}(\mathbf{x})]] \tag{4.12}$$

$$= \mathbb{E}_{\mathbf{x}}\big[\mathbf{x}\mathbf{k}^\top(\mathbf{x})\big]\mathbf{K}^{-1}\mathbf{Y} - \mathbf{m}\boldsymbol{\mu}_{\mathrm{G}}^\top \tag{4.13}$$

Summary of the proposed GPQ moment transform is given below.

General GPQ moment transform
The general GPQ based Gaussian approximation to the joint distribution of $\mathbf{x}$ and a transformed random variable $\mathbf{y} = \mathbf{g}(\mathbf{x})$, where $\mathbf{x} \sim \mathrm{N}(\mathbf{m}, \mathbf{P})$, is given by

$$\begin{bmatrix}\mathbf{x}\\ \mathbf{y}\end{bmatrix} \sim \mathrm{N}\left(\begin{bmatrix}\mathbf{m}\\ \boldsymbol{\mu}_{\mathrm{G}}\end{bmatrix}, \begin{bmatrix}\mathbf{P} & \mathbf{C}_{\mathrm{G}}\\ \mathbf{C}_{\mathrm{G}}^\top & \boldsymbol{\Pi}_{\mathrm{G}}\end{bmatrix}\right) \tag{4.14}$$

where the transformed moments are computed as

$$\boldsymbol{\mu}_{\mathrm{G}} = \mathbf{Y}^\top\mathbf{w}, \tag{4.15}$$

$$\boldsymbol{\Pi}_{\mathrm{G}} = \mathbf{Y}^\top\mathbf{W}\mathbf{Y} - \boldsymbol{\mu}\boldsymbol{\mu}^\top + \sigma^2\mathbf{I}, \tag{4.16}$$

$$\mathbf{C}_{\mathrm{G}} = \mathbf{W}_c\mathbf{Y} - \mathbf{m}\boldsymbol{\mu}^\top, \tag{4.17}$$

$$\sigma^2 = \bar{k} - \mathrm{tr}\left(\mathbf{Q}\mathbf{K}^{-1}\right), \tag{4.18}$$

and where $\big[\mathbf{Y}\big]_{*e} = \begin{bmatrix}y_1^e & \dots & y_N^e\end{bmatrix}^\top$ are the function values of the e-th output dimension of $\mathbf{g}(\mathbf{x})$. The kernel matrix $\mathbf{K}$ is defined in section 3.1 and the GPQ weights are

$$\mathbf{w} = \mathbf{K}^{-1}\mathbf{q},\ \mathbf{W} = \mathbf{K}^{-1}\mathbf{Q}\mathbf{K}^{-1} \text{ and } \mathbf{W}_c = \mathbf{R}\mathbf{K}^{-1}, \tag{4.19}$$

where

$$[\mathbf{q}]_i = \mathbb{E}_{\mathbf{x}}[k(\mathbf{x},\mathbf{x}_i)], \tag{4.20}$$

$$[\mathbf{Q}]_{ij} = \mathbb{E}_{\mathbf{x}}[k(\mathbf{x},\mathbf{x}_i)k(\mathbf{x},\mathbf{x}_j)], \tag{4.21}$$

$$[\mathbf{R}]_{*j} = \mathbb{E}_{\mathbf{x}}[\mathbf{x}k(\mathbf{x},\mathbf{x}_j)], \tag{4.22}$$

$$\bar{k} = \mathbb{E}_{\mathbf{x}}[k(\mathbf{x},\mathbf{x})]. \tag{4.23}$$

The sigma-points $\mathbf{x}_i$ can be chosen *arbitrarily*.

4.2 Properties of the general GPQ transform

An important requirement of moment transforms is that they produce valid covariance matrices. 4.2 given below states that the proposed GPQ transform always produces positive semi-definite covariance matrix. For proof we use the following lemma.

Lemma 4.1. *For any $m \times n$ matrix $\mathbf{X}$ and positive definite $n \times n$ matrix $\mathbf{A}$, the matrix $\mathbf{X}\mathbf{A}\mathbf{X}^\top$ is positive semi-definite.*

Proof. See [Horn and Johnson, 1990, Observation 7.1.6, p. 399]. □

In the following, let $\mathbf{A} \succeq 0 \ \Leftrightarrow\ \mathbf{x}^\top\mathbf{A}\mathbf{x} \geq 0,\ \forall \mathbf{x} \in \mathbb{R}^n$ for any $n \times n$ matrix $\mathbf{A}$.

Theorem 4.2. *The GPQ transformed covariance is positive semi-definite.*

Proof. Using the expressions for the GPQ weights from eq. (4.19), we can write

$$\boldsymbol{\Pi} = \mathbf{Y}^\top \mathbf{K}^{-1}\big(\mathbf{Q} - \mathbf{q}\mathbf{q}^\top\big)\mathbf{K}^{-1}\mathbf{Y} + \sigma^2\mathbf{I} = \mathbf{Z}^\top\widetilde{\mathbf{Q}}\mathbf{Z} + \sigma^2\mathbf{I} = \widetilde{\boldsymbol{\Pi}} + \sigma^2\mathbf{I},$$

where $\widetilde{\boldsymbol{\Pi}} = \mathbf{Z}^\top\widetilde{\mathbf{Q}}\mathbf{Z}$, $\mathbf{Z} = \mathbf{K}^{-1}\mathbf{Y}$ and $\widetilde{\mathbf{Q}} = \mathbf{Q} - \mathbf{q}\mathbf{q}^\top$. We recognize that $\widetilde{\mathbf{Q}} = \mathbb{C}[\mathbf{k}(\mathbf{x})] = \mathbb{E}\big[\mathbf{k}(\mathbf{x})\mathbf{k}^\top(\mathbf{x})\big] - \mathbb{E}[\mathbf{k}(\mathbf{x})]\mathbb{E}[\mathbf{k}(\mathbf{x})]^\top$. From the property of covariance matrices it follows that $\widetilde{\mathbf{Q}} \succeq 0$. The 4.1 implies that $\widetilde{\boldsymbol{\Pi}} = \mathbf{Z}^\top\widetilde{\mathbf{Q}}\mathbf{Z} \succeq 0$ for any matrix $\mathbf{Z}$. Finally, since $\sigma^2 \geq 0$, we have that $\boldsymbol{\Pi} = \widetilde{\boldsymbol{\Pi}} + \sigma^2\mathbf{I} \succeq 0$. □

Evidently, the GPQ moment transform hinges upon the kernel expectations given by eqs. (4.20) to (4.22). Since we are already using one quadrature to approximate moments, it is thus preferable that these expectations be analytically tractable. A list of tractable kernel-density pairs is provided in [Briol et al., 2015]. A popular choice in many applications is an RBF (Gaussian) kernel, expectations of which are summarized below.

Theorem 4.3 (GPQ transform with RBF kernel)**.** *Assuming a change of variables has taken place in the Gaussian weighted integrals given by eq.* (2.6) *and the kernel is of the form*

$$k(\boldsymbol{\xi}, \boldsymbol{\xi}') = \alpha^2 \exp\big(-\tfrac{1}{2}(\boldsymbol{\xi} - \boldsymbol{\xi}')^\top \boldsymbol{\Lambda}^{-1}(\boldsymbol{\xi} - \boldsymbol{\xi}')\big), \tag{4.24}$$

where α *is scaling parameter and* $\boldsymbol{\Lambda} = \operatorname{diag}\big(\begin{bmatrix}\ell_1^2 & \dots & \ell_D^2\end{bmatrix}\big)$ *is lengthscale, then the expectations given by eqs.* (4.20) *to* (4.22) *take on the form*

$$[\mathbf{q}]_i = \alpha^2\big|\boldsymbol{\Lambda}^{-1} + \mathbf{I}\big|^{-\frac{1}{2}} \exp\big(-\tfrac{1}{2}\boldsymbol{\xi}_i^\top(\boldsymbol{\Lambda} + \mathbf{I})^{-1}\boldsymbol{\xi}_i\big), \tag{4.25}$$

$$[\mathbf{Q}]_{ij} = \alpha^4\big|2\boldsymbol{\Lambda}^{-1} + \mathbf{I}\big|^{-\frac{1}{2}} \exp\Big(-\tfrac{1}{2}\Big(\boldsymbol{\xi}_i^\top\boldsymbol{\Lambda}^{-1}\boldsymbol{\xi}_i + \boldsymbol{\xi}_j\boldsymbol{\Lambda}^{-1}\boldsymbol{\xi}_j - \mathbf{z}_{ij}^\top\big(2\boldsymbol{\Lambda}^{-1} + \mathbf{I}\big)^{-1}\mathbf{z}_{ij}\Big)\Big), \tag{4.26}$$

$$[\mathbf{R}]_{*j} = \alpha^2\big|\boldsymbol{\Lambda}^{-1} + \mathbf{I}\big|^{-\frac{1}{2}} \exp\big(-\tfrac{1}{2}\boldsymbol{\xi}_j^\top(\boldsymbol{\Lambda} + \mathbf{I})^{-1}\boldsymbol{\xi}_j\big)(\boldsymbol{\Lambda} + \mathbf{I})^{-1}\boldsymbol{\xi}_j, \tag{4.27}$$

$$\bar{k} = \alpha^2 \tag{4.28}$$

where $\mathbf{z}_{ij} = \boldsymbol{\Lambda}^{-1}(\boldsymbol{\xi}_i + \boldsymbol{\xi}_j)$.

Proof. Expressions can be derived by writing the RBF kernel as a Gaussian and making use of the formulas for the product of two Gaussian densities (and the normalizing constant). For space reasons, we omit the fairly straightforward, but nevertheless lengthy and tedious, derivations. □

It is now evident that the GPQ transformed moments depend on the kernel parameters which need to be set prior to computing the weights. The form of $\boldsymbol{\Lambda}$ in the RBF kernel formulation above exhibits, so called, automatic relevance determination (ARD). That is to say, by optimizing the lengthscales ℓ_d dimensions contributing most to the variability in the data can be discovered, where a small ℓ_d would indicate high relevance of the d-th dimension. A typical approach in GP regression would be to optimize the kernel parameters by marginal likelihood (evidence) maximization. However, in the BQ setting this method would likely yield unreliable parameter estimates due to the inherently minimal amount of data available. For these reasons, we resorted to a manual choice of the parameter values which were mostly informed by the prior knowledge of the integrated function. In the following theorem, we prove the weight independence on the kernel scaling parameter.

Theorem 4.4 (Kernel scaling independence)**.** *Assume a scaled version of a kernel is used, so that* $\bar{k}(\mathbf{x}, \mathbf{x}') = c \cdot k(\mathbf{x}, \mathbf{x}')$*, then the weights of the GPQ transform given in eq.* (4.19) *are independent of the scaling parameter* c*.*

Proof. Define a scaled kernel matrix $\mathbf{K}' = c\mathbf{K}$, and scaled kernel expectations $[\mathbf{q}']_i = c\mathbb{E}_{\mathbf{x}}[k(\mathbf{x}, \mathbf{x}_i)]$, $[\mathbf{Q}']_{ij} = c^2\mathbb{E}_{\mathbf{x}}[k(\mathbf{x}, \mathbf{x}_i)k(\mathbf{x}, \mathbf{x}_j)]$, $[\mathbf{R}']_{*j} = c\mathbb{E}_{\mathbf{x}}[\mathbf{x}k(\mathbf{x}, \mathbf{x}_j)]$. Plugging into the expressions for the GPQ weights from the eq. (4.19), we get

$$\mathbf{w}' = \mathbf{q}'\big(\mathbf{K}'\big)^{-1} = cc^{-1}\mathbf{q}\mathbf{K}^{-1} = \mathbf{w}, \tag{4.29}$$

$$\mathbf{W}' = \big(\mathbf{K}'\big)^{-1}\mathbf{Q}'\big(\mathbf{K}'\big)^{-1} = c^2c^{-2}\mathbf{K}^{-1}\mathbf{Q}\mathbf{K}^{-1} = \mathbf{W}, \tag{4.30}$$

$$\mathbf{W}_c' = \mathbf{R}'\big(\mathbf{K}'\big)^{-1} = cc^{-1}\mathbf{R}\mathbf{K}^{-1} = \mathbf{W}_c. \tag{4.31}$$

□

Corollary 4.5. *The kernel scaling affects only the additive term in the transformed covariance, which becomes* $\sigma^2 = c\big[\bar{k} - \operatorname{tr}\big(\mathbf{Q}\mathbf{K}^{-1}\big)\big]$. *The transformed mean and cross-covariance are unaffected by the scaling.*

4.3 Moment transforms in filtering

An important application area for the moment transforms is in local filtering algorithms, which estimate an evolving system state from noisy measurements. Filters are essentially inference algorithms operating on a state-space model which is given by the two discrete-time equations

$$\mathbf{x}_k = \mathbf{f}(\mathbf{x}_{k-1}) + \mathbf{q}_k, \tag{4.32}$$

$$\mathbf{z}_k = \mathbf{h}(\mathbf{x}_k) + \mathbf{r}_k, \tag{4.33}$$

describing the system dynamics and the state observation (measurement) process respectively. The evolution of the system state $\mathbf{x}_k$ is described by eq. (4.32) where $\mathbf{f} : \mathbb{R}^D \to \mathbb{R}^D$ is the dynamics and the $\mathbf{q}_k \sim \mathrm{N}(\mathbf{0}, \mathbf{Q})$ is the state noise. The measurements $\mathbf{z}_k$ are produced by mapping the state through the observation model $\mathbf{h} : \mathbb{R}^D \to \mathbb{R}^E$ and adding the measurement noise $\mathbf{r}_k \sim \mathrm{N}(\mathbf{0}, \mathbf{R})$. Typically, $E \leq D$. Let $\mathbf{z}_{1:k} = \{\mathbf{z}_1, \ldots, \mathbf{z}_k\}$ denote a set of measurements up to time step k. From a probabilistic standpoint, the filtering problem is about inferring a posterior distribution

$$p(\mathbf{x}_k \mid \mathbf{z}_{1:k}) \propto p(\mathbf{x}_k, \mathbf{z}_k \mid \mathbf{z}_{1:k-1}) = p(\mathbf{z}_k \mid \mathbf{x}_k) p(\mathbf{x}_k \mid \mathbf{z}_{1:k-1}) \tag{4.34}$$

where the likelihood $p(\mathbf{z}_k \mid \mathbf{x}_k)$ is obtained from eq. (4.33) and the prior $p(\mathbf{x}_k \mid \mathbf{z}_{1:k-1})$ is given by the Chapman-Kolmogorov equation

$$p(\mathbf{x}_k \mid \mathbf{z}_{1:k-1}) = \int p(\mathbf{x}_k \mid \mathbf{x}_{k-1}) p(\mathbf{x}_{k-1} \mid \mathbf{z}_{1:k-1}) \, \mathrm{d}\mathbf{x}_{k-1} \tag{4.35}$$

where the transition density $p(\mathbf{x}_k \mid \mathbf{x}_{k-1})$ is obtained from the system dynamics given by eq. (4.32). Local Gaussian filters make the simplifying assumption that the state and measurement are jointly Gaussian distributed, so that

$$p(\mathbf{x}_k, \mathbf{z}_k \mid \mathbf{z}_{1:k-1}) \approx \mathrm{N}\left(\begin{bmatrix} \mathbf{x}_k \\ \mathbf{z}_k \end{bmatrix} \middle| \begin{bmatrix} \mathbf{m}^x_{k|k-1} \\ \mathbf{m}^z_{k|k-1} \end{bmatrix}, \begin{bmatrix} \mathbf{P}^x_{k|k-1} & \mathbf{P}^{xz}_{k|k-1} \\ \mathbf{P}^{xz}_{k|k-1} & \mathbf{P}^z_{k|k-1} \end{bmatrix} \right), \tag{4.36}$$

where the index notation $k|k-1$ means that the relevant quantity at time k is computed from $\mathbf{z}_{1:k-1}$. Advantage of this simplification is that the posterior is now parametrized by the conditional mean and covariance, which are available in closed form. Using an update rule the state and measurement moments of the joint in eq. (4.36)

$$\mathbf{m}^x_{k|k} = \mathbf{m}^x_{k|k-1} + \mathbf{P}^{xz}_{k|k-1} (\mathbf{P}^z_{k|k-1})^{-1} \big(\mathbf{z}_k - \mathbf{m}^z_{k|k-1}\big), \tag{4.37}$$

$$\mathbf{P}^x_{k|k} = \mathbf{P}^x_{k|k-1} - \mathbf{P}^{xz}_{k|k-1} (\mathbf{P}^z_{k|k-1})^{-1} (\mathbf{P}^{xz}_{k|k-1})^\top, \tag{4.38}$$

are combined to arrive at the approximate conditional mean and covariance of the state.

To compute state predictive moments $\mathbf{m}^x_{k|k-1}$, $\mathbf{P}^x_{k|k-1}$ a moment transform is applied in a setting where the input moments are $\mathbf{m}^x_{k-1|k-1}$, $\mathbf{P}^x_{k-1|k-1}$ and the nonlinearity is the dynamics $\mathbf{f}(\mathbf{x}_{k-1})$. Similarly, the measurement moments $\mathbf{m}^z_{k|k-1}$, $\mathbf{P}^z_{k|k-1}$, $\mathbf{P}^{xz}_{k|k-1}$ are obtained by applying a moment transform on input moments $\mathbf{m}^x_{k|k-1}$, $\mathbf{P}^x_{k|k-1}$ with nonlinearity $\mathbf{h}(\mathbf{x}_k)$.

5 Experiments

The proposed GPQ moment transform is first tested on a polar-to-Cartesian coordinate transformation while the later experiments focus on applications in nonlinear filtering. In all cases the GPQ transform uses the RBF kernel given by eq. (4.24). Since the sigma-point locations are not prescribed and their choice is entirely arbitrary, we used the point sets of the classical rules mentioned in Section 2 for all examples. The acronym GPQKF denotes all nonlinear Kalman filters based on the GPQ regardless of which point set they use.

5.1 Mapping from Polar to Cartesian Coordinates

Conversion from polar to Cartesian coordinates is a ubiquitous nonlinearity appearing in radar sensors or laser range finders and is given by

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} r\cos(\theta) \\ r\sin(\theta) \end{bmatrix}. \tag{5.1}$$

Since the mapping is conditionally linear (for fixed θ) and we use a kernel with ARD in the moment transform, we can exploit this fact and set the kernel lengthscales to $\ell = \begin{bmatrix} 60 & 6 \end{bmatrix}$ while the scaling was set to $\alpha = 1$. Note, that we set the lengthscale corresponding to range to a relatively large value. This is because the larger lengthscales in the kernel correspond to a slower variation in the approximated function.

We compared the performance of the spherical radial transform (SR), which is basis of the cubature Kalman filter [Arasaratnam and Haykin, 2009], and the GPQ transform with SR points (GPQ-SR) for 100 different input moments. The 10 different positions on a spiral in polar coordinates were chosen as input means $\mathbf{m}_i = \begin{bmatrix} r_i & \theta_i \end{bmatrix}$. For each mean we assigned 10 different input covariance matrices $\mathbf{P}_j = \operatorname{diag}\big(\begin{bmatrix} \sigma_r^2 & \sigma_{\theta,j}^2 \end{bmatrix} \big)$, where $\sigma_r = 0.5\,\mathrm{m}$ and $\sigma_{\theta,j} \in [6^\circ, 36^\circ]$ for $j = 1, \ldots, 10$. Figure 3 depicts the input means in polar coordinates. As a measure of an

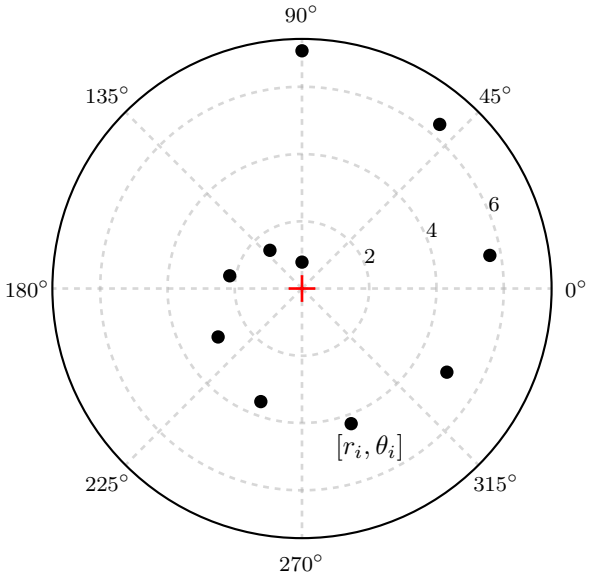

Figure 3: Input means are placed on a spiral. For each input mean $\mathbf{m}_i = \begin{bmatrix} r_i & \theta_i \end{bmatrix}$ (black dot) the radius variance is fixed at $\sigma_r = 0.5\,\mathrm{m}$ and 10 azimuth variances are considered so that $\sigma_\theta \in [6^\circ, 36^\circ]$.

agreement between two Gaussian densities we used the symmetrized KL-divergence given by

$$\begin{aligned} \mathrm{SKL} &= \frac{1}{2}\{\mathbb{KL}(\mathrm{N}(\mathbf{y} \,|\, \boldsymbol{\mu}, \boldsymbol{\Pi}) \,\|\, \mathrm{N}(\mathbf{y} \,|\, \boldsymbol{\mu}_\mathrm{G}, \boldsymbol{\Pi}_\mathrm{G})) + \mathbb{KL}(\mathrm{N}(\mathbf{y} \,|\, \boldsymbol{\mu}_\mathrm{G}, \boldsymbol{\Pi}_\mathrm{G}) \,\|\, \mathrm{N}(\mathbf{y} \,|\, \boldsymbol{\mu}, \boldsymbol{\Pi}))\} && (5.2) \\ &= \frac{1}{4}\Big\{ (\boldsymbol{\mu} - \boldsymbol{\mu}_\mathrm{G})^\top \boldsymbol{\Pi}^{-1} (\boldsymbol{\mu} - \boldsymbol{\mu}_\mathrm{G}) + (\boldsymbol{\mu}_\mathrm{G} - \boldsymbol{\mu})^\top \boldsymbol{\Pi}_\mathrm{G}^{-1} (\boldsymbol{\mu}_\mathrm{G} - \boldsymbol{\mu}) + \operatorname{tr}(\boldsymbol{\Pi}^{-1}\boldsymbol{\Pi}_\mathrm{G}) + \operatorname{tr}(\boldsymbol{\Pi}_\mathrm{G}^{-1}\boldsymbol{\Pi}) - 2E \Big\}, && (5.3) \end{aligned}$$

where $E = \dim(\mathbf{y})$. The ground truth transformed mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Pi}$ were computed using the Monte Carlo method with 10 000 samples. Two SKL scores were considered; the average over means and an average over azimuth variances.

The Figure 4 shows the SKL score calculated for each configuration on the spiral. The left pane of Figure 4 shows results for individual means averaged over the azimuth variances, whereas the right pane displays averaged SKL over the means. In both cases our proposed moment transform outperforms the classical quadrature transform with the same SR point set.

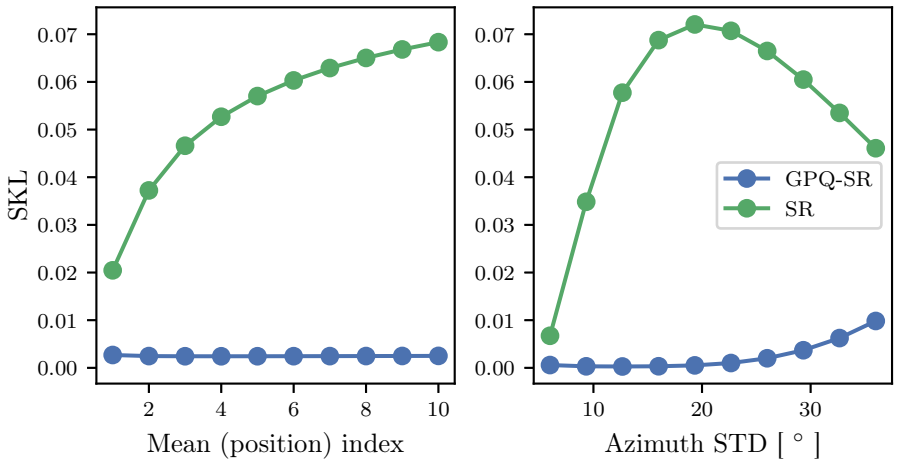

Figure 4: Performance comparison of the spherical radial (SR) and GPQ with SR points (GPQ-SR) moment transforms in terms of averaged symmetrized KL-divergence. Left: average over a range of azimuth variances; Right: average over the range of input means (positions on the spiral).

5.2 UNGM

The performance of nonlinear sigma-point filters based on GPQ transform was first tested in [Prüher and Šimandl, 2016] on a univariate non-stationary growth model (UNGM), where the system dynamics and the observation model are given by

$$x_k = \frac{1}{2}x_{k-1} + \frac{25x_{k-1}}{1+x_{k-1}^2} + 8\cos(1.2\,k) + q_{k-1}\,, \tag{5.4}$$

$$z_k = \frac{1}{20}x_{k-1}^2 + r_k\,, \tag{5.5}$$

with the state noise $q_{k-1} \sim \mathcal{N}(0, 10)$, measurement noise $r_k \sim \mathcal{N}(0, 1)$ and initial conditions $x_{0|0} \sim \mathcal{N}(0, 5)$. This model is frequently used as a benchmark in nonlinear filtering [Gordon et al., 1993; Kitagawa, 1996].

For this problem, all of the considered GPQKFs used the same kernel scaling $\alpha = 1$. The lengthscale was set to $\ell = 3.0$ for the UT, $\ell = 0.3$ for SR and GH-5, and $\ell = 0.1$ for all higher-order GH sigma-point sets. The GPQKFs that used UT and GH sigma-points of order 5, 7, 10, 15 and 20 were compared with their classical quadrature-based counterparts, namely, the UKF and the GHKF of the same orders. The UKF operated with $\kappa = 0$. We performed 100 simulations, each for $K = 500$ time steps.

For evaluation of the filter performance, we used the root-mean-square error (RMSE)

$$\mathrm{RMSE} = \sqrt{\frac{1}{K}\sum_{k=1}^{K}\left(\mathbf{x}_k - \mathbf{m}_{k|k}^{\mathbf{x}}\right)^2} \tag{5.6}$$

to measure the overall difference between the state estimate $\mathbf{m}_{k|k}^{\mathbf{x}}$ and the true state $\mathbf{x}_k$ across all time steps. The negative log-likelihood of the state estimate $\mathbf{m}_{k|k}^{\mathbf{x}}$ and covariance $\mathbf{P}_k^{\mathbf{x}}$

$$\mathrm{NLL} = -\log p(\mathbf{x}_k \mid \mathbf{z}_{1:k}) = \frac{1}{2}\Big[\log|2\pi\mathbf{P}_k^{\mathbf{x}}| + (\mathbf{x}_k - \mathbf{m}_{k|k}^{\mathbf{x}})^\top (\mathbf{P}_k^{\mathbf{x}})^{-1}(\mathbf{x}_k - \mathbf{m}_{k|k}^{\mathbf{x}})\Big] \tag{5.7}$$

was used to measure the overall model fit [Gelman, 2013]. As a metric that takes into account the estimated state covariance, the inclination indicator [Li and Zhao, 2006] given by

$$\nu = \frac{10}{K}\sum_{k=1}^{K}\log_{10}\frac{\left(\mathbf{x}_k - \mathbf{m}_{k|k}^{\mathbf{x}}\right)^\top \left(\mathbf{P}_{k|k}^{\mathbf{x}}\right)^{-1}\left(\mathbf{x}_k - \mathbf{m}_{k|k}^{\mathbf{x}}\right)}{\left(\mathbf{x}_k - \mathbf{m}_{k|k}^{\mathbf{x}}\right)^\top \boldsymbol{\Sigma}_k^{-1}\left(\mathbf{x}_k - \mathbf{m}_{k|k}^{\mathbf{x}}\right)} \tag{5.8}$$

Point set	N	GPQ	Classical
SR	2	6.157 ± 0.071	13.652 ± 0.253
UT	3	7.124 ± 0.131	7.103 ± 0.130
GH5	5	8.371 ± 0.128	10.466 ± 0.198
GH7	7	8.360 ± 0.043	9.919 ± 0.215
GH10	10	7.082 ± 0.038	8.035 ± 0.193
GH15	15	6.944 ± 0.048	8.224 ± 0.188
GH20	20	6.601 ± 0.058	7.406 ± 0.193

Table 1: The average root-mean-square error.

Point set	N	GPQ	Classical
SR	2	3.328 ± 0.026	56.570 ± 2.728
UT	3	4.970 ± 0.343	5.306 ± 0.481
GH5	5	4.088 ± 0.064	14.722 ± 0.829
GH7	7	4.045 ± 0.017	12.395 ± 0.855
GH10	10	3.530 ± 0.012	7.565 ± 0.534
GH15	15	3.468 ± 0.014	7.142 ± 0.557
GH20	20	3.378 ± 0.017	5.664 ± 0.488

Table 2: The average negative log-likelihood.

was used, where $\mathbf{\Sigma}_k$ is the sample mean-square-error matrix. The filter is said to be optimistic if it underestimates the actual error, which is indicated by $\nu > 0$ and vice versa. A perfectly credible filter would provide $\nu = 0$, that is, it would neither overestimate nor underestimate the actual error.

Tables 1 to 3 show average values of the performance criteria across simulations with bootstrapped estimates of ± 2 standard deviations [Wasserman, 2007]. As evidenced by the results in Table 1, the Bayesian quadrature achieves superior RMSE performance for all sigma-point sets. In the classical quadrature case the performance improves with increasing number of sigma-points used. Table 2 shows that the performance of GPQKF is clearly superior in terms of NLL, which indicates that the estimates produced by the GPQ-based filters are better representations of the unknown true state development. The self-assessment of the filter performance is more credible in the case of GPQ, as indicated by lower inclination ν in the Table 3. This indicates that the GPQ-based filters are more conservative in their covariance estimates - a consequence of including additional uncertainty (integral variance), which the classical quadrature-based filters do not employ. Also note, that the variance of all the evaluated criteria for GPQ-based filters is mostly an order of magnitude lower.

Point set	N	GPQ	Classical
SR	2	1.265 ± 0.010	18.585 ± 0.045
UT	3	0.363 ± 0.108	0.897 ± 0.088
GH5	5	4.549 ± 0.013	9.679 ± 0.068
GH7	7	4.638 ± 0.006	8.409 ± 0.076
GH10	10	2.520 ± 0.006	5.315 ± 0.058
GH15	15	2.331 ± 0.008	5.424 ± 0.059
GH20	20	1.654 ± 0.007	4.105 ± 0.055

Table 3: The average inclination indicator.

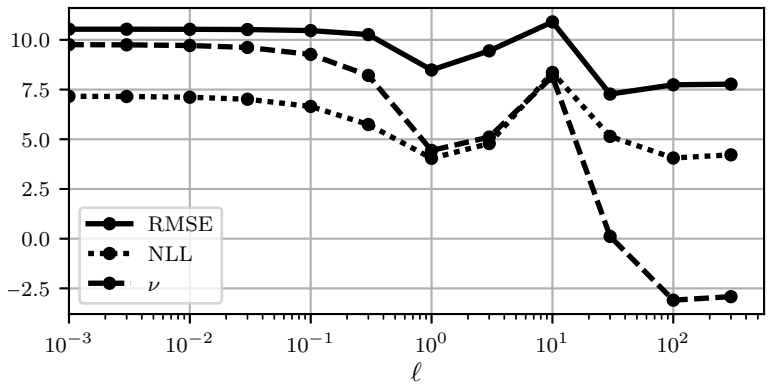

Figure 5: Sensitivity of GPQKF performance (using UT sigma-points) to changes in the lengthscale parameter ℓ. The choice $\ell = 3$ minimizes RMSE and yields nearly optimal inclination ν.

To achieve competitive results, the kernel lengthscale ℓ had to be manually set for each filter separately. This was done by running the filters with increasing lengthscale, plotting the performance metrics and choosing the value which gave the smallest RMSE and the inclination closest to zero. Figure 5 illustrates the effect of changing lengthscale on the overall performance of the GPQKF with UT sigma-points.

5.3 Target Tracking

As a more application oriented example, we considered a target tracking scenario adopted from [Athans et al., 1968; Julier et al., 2000]. A spherical object falls down from high altitude entering the Earth's atmosphere with high velocity. The nonlinear dynamics is described by the following set of differential equations

$$\dot{p}(t) = -v(t) + q_1(t), \tag{5.9}$$

$$\dot{v}(t) = -v^2(t)\theta(t)e^{-\gamma p(t)} + q_2(t), \tag{5.10}$$

$$\dot{\theta}(t) = q_3(t), \tag{5.11}$$

where $\gamma = 0.164$ is a constant and the system state $\mathbf{x} = \begin{bmatrix} p & v & \theta \end{bmatrix}$ consists of position (altitude) p, velocity v and a constant ballistic parameter θ. The zero-mean state noise is characterized by $\mathbb{E}[q_i(t)q_j(s)] = \mathbf{Q}\delta(t-s)\delta(i-j)$, where $\mathbf{q} = \begin{bmatrix} q_1 & q_2 & q_3 \end{bmatrix}$. The range measurements are produced at discrete time intervals by a radar positioned at the altitude of 30 km and 30 km horizontally to the vertical path of the falling object. Thus the observation model is

$$y_k = \sqrt{s_x^2 + (s_y - p_k)^2} + r_k, \tag{5.12}$$

where $(s_x,\ s_y)$ is the radar position. The measurements were generated with frequency 10 Hz and the measurement noise is zero-mean with variance $\sigma_y^2 = 9.2903 \times 10^{-4}\,\mathrm{km}^2$. The mean and covariance of the system initial condition were set to

$$\mathbf{x}_0 = \begin{bmatrix} 90\,\mathrm{km} & 6\,\mathrm{km\,s^{-1}} & 1.5 \end{bmatrix} \tag{5.13}$$

$$\mathbf{P}_0 = \mathrm{diag}\big(\begin{bmatrix} 0.0929\,\mathrm{km}^2 & 1.4865\,\mathrm{km}^2\,\mathrm{s}^{-2} & 10^{-4} \end{bmatrix} \big) \tag{5.14}$$

while the filter used different initial state estimate

$$\mathbf{m}_{0|0}^x = \begin{bmatrix} 90\,\mathrm{km} & 6\,\mathrm{km\,s^{-1}} & 1.7 \end{bmatrix} \tag{5.15}$$

$$\mathbf{P}_{0|0}^x = \mathrm{diag}\big(\begin{bmatrix} 0.0929\,\mathrm{km}^2 & 1.4865\,\mathrm{km}^2\,\mathrm{s}^{-2} & 10 \end{bmatrix} \big) \tag{5.16}$$

which implies a lighter object than in reality.

In the experiments, we focused on the comparison of our GPQKF with the UT points and the UKF, because this filter was previously used in [Julier et al., 2000] to demonstrate its superiority over the EKF on the same tracking problem. The parameters of the UKF were set to $\kappa = 0$, $\alpha = 1$, $\beta = 2$ following the recommended heuristics [Särkkä, 2013]. The GPQKF used different kernel parameters for the dynamics, $\alpha_f = 0.5$, $\ell_f = \begin{bmatrix}10 & 10 & 10\end{bmatrix}$, and the measurement nonlinearity, $\alpha_h = 0.5$, $\ell_h = \begin{bmatrix}15 & 20 & 20\end{bmatrix}$. All filters operated with a discrete-time model obtained by Euler approximation with step size $\Delta t = 0.1\,\mathrm{s}$. The discretized model is given by

$$p(k+1) = p(k) - \Delta t\, v(k) + q_1(k), \tag{5.17}$$

$$v(k+1) = v(k) - \Delta t\, v^2(k)\theta(k)\exp(-\gamma p(k)) + q_2(k), \tag{5.18}$$

$$\theta(k+1) = \theta(k) + q_3(k). \tag{5.19}$$

We generated 100 truth trajectories by simulating the continuous-time dynamics, given by the eqs. (5.9) to (5.11), for 30 time steps by 4th-order Runge-Kutta scheme and computed the average RMSE and inclination indicator ν for both tested filters. Figure 6 shows realizations of the altitude and velocity trajectories along with the average trajectory. Note that when the object is passing directly in front of the radar at approximately $t = 10$ s (i.e. altitude 30 km), the system is almost unobservable.

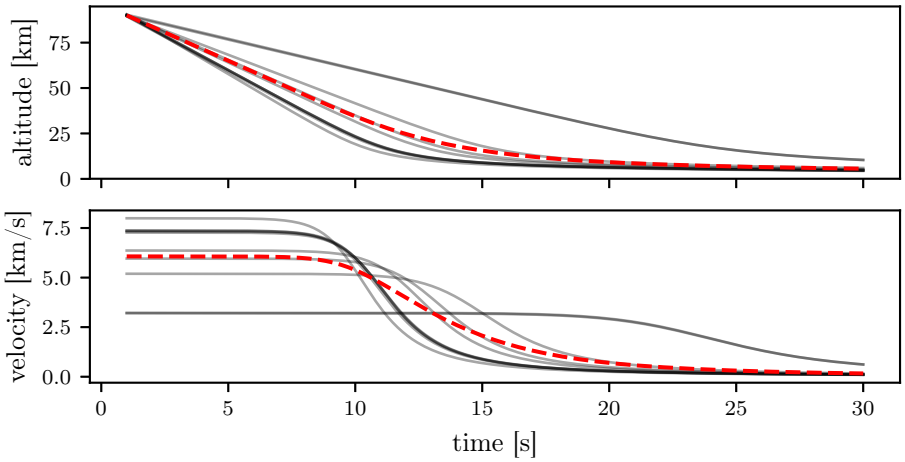

Figure 6: Altitude and velocity evolution in time. Trajectory realizations (black) and the average trajectory (red). The greatest deceleration occurs in the period from 10 to 20 seconds.

Figure 7 depicts the RMSE for each time step averaged over trajectory simulations. The RMSE of the GPQKF tends to be better for all state vector components. The biggest difference is evident in the RMSE of the ballistic parameter where GPQKF shows significantly better performance during the period of the greatest deceleration. Overall, the UKF shows signs of an unbalanced estimator as evidenced from Figure 8, where the inclination ν rises significantly above zero, indicating excessive optimism. The GPQKF manages to stay mostly balanced (ν wobbles around zero) with the exception of velocity, where it tips toward pessimism towards the end of the trajectory. This behaviour is mostly likely caused by the inclusion of additional functional uncertainty in the transformed covariance as shown in eqs. (4.5) and (4.6).

6 Conclusion

In this article, we have shown how a Bayesian view of quadrature can be leveraged for the design of general purpose moment transform. Unlike the classical transforms, the proposed GPQ transform accounts for the integration error incurred in computing the mean by inflating the transformed covariance. The underlying model in the proposed transform is a non-parametric GP regression model, which brings a lot of advantages.

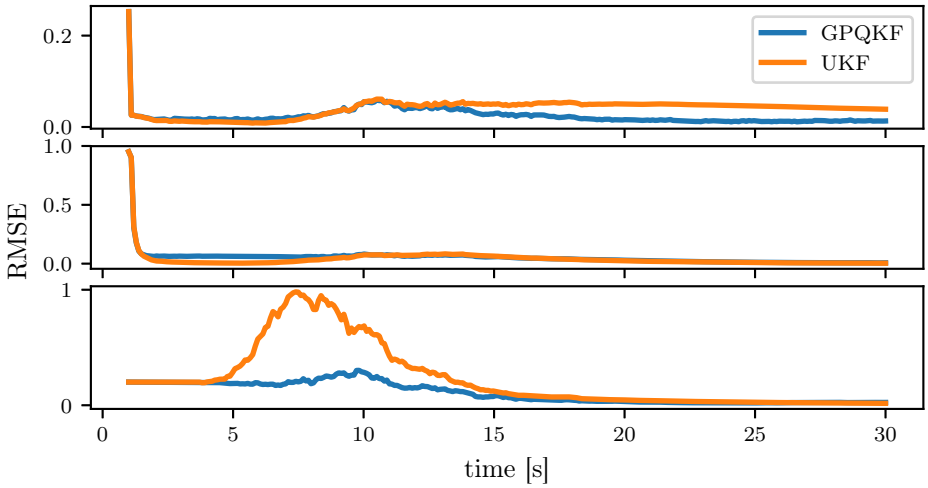

Figure 7: Evolution of the average RMSE in time for the GPQKF with the UT points and the UKF. From top to bottom: position, velocity and ballistic parameter.

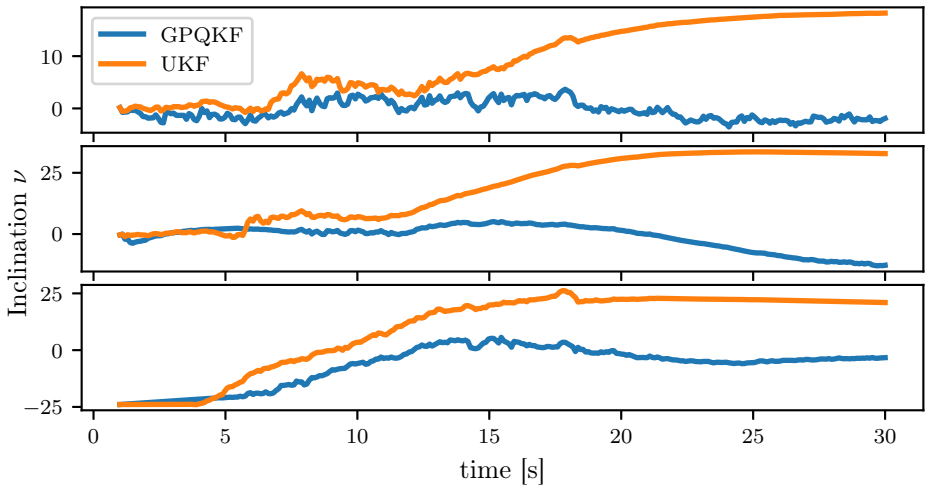

Figure 8: Evolution of the average inclination in time for the GPQKF with the UT points and the UKF. From top to bottom: position, velocity and ballistic parameter.

Namely, the transform is not restricted by polynomial assumptions on the integrand (unlike the classical methods) and it quantifies predictive uncertainty, which eventually translates into integral uncertainty. The proposed moment transform is entirely general, in that the equations hold for any kernel and input density, however, analytically tractable kernel-density pairs are preferable. We showed that the transform may outperform classical transforms on a coordinate conversion and two nonlinear sigma-point filtering examples. In both experiments, the filters based on the GPQ give more realistic estimates of the covariance, hence are better at self-assessing their estimation error. Currently, the biggest challenge is finding optimal values of the kernel parameters.

References

I. Arasaratnam and S. Haykin. Cubature Kalman Filters. *IEEE Transactions on Automatic Control*, 54(6): 1254–1269, 2009.

M. Athans, R. Wishner, and A. Bertolini. Suboptimal State Estimation for Continuous-Time Nonlinear Systems

from Discrete Noisy Measurements. *IEEE Transactions on Automatic Control*, 13(5):504–514, 1968. doi: 10.1109/TAC.1968.1098986.

R. Bhar. *Stochastic filtering with applications in finance*. World Scientific, 2010. ISBN 978-981-4304-85-6.

F.-X. Briol, C. J. Oates, M. Girolami, M. A. Osborne, and D. Sejdinovic. Probabilistic Integration: A Role for Statisticians in Numerical Analysis? *ArXiv e-prints*, December 2015.

M. P. Deisenroth, M. F. Huber, and U. D. Hanebeck. Analytic moment-based Gaussian process filtering. In *Proceedings of the 26th Annual International Conference on Machine Learning - ICML '09*, pages 1–8. ACM Press, 2009. doi: 10.1145/1553374.1553403.

M. P. Deisenroth, R. D. Turner, M. F. Huber, U. D. Hanebeck, and C. E. Rasmussen. Robust Filtering and Smoothing with Gaussian Processes. *IEEE Transactions on Automatic Control*, 57(7):1865–1871, 2012. doi: 10.1109/TAC.2011.2179426.

P. Diaconis. Bayesian numerical analysis. In S. S. Gupta and J. O. Berger, editors, *Statistical Decision Theory and Related Topics IV*, pages 163–175. Springer, 1988.

J. Duník, O. Straka, and M. Šimandl. Stochastic Integration Filter. *IEEE Transactions on Automatic Control*, 58(6):1561–1566, 2013.

W. Gautschi. *Orthogonal Polynomials: Computation and Approximation*. Numerical Mathematics and Scientific Computation. Oxford University Press, 2004. ISBN 978-0198506720.

A. Gelman. *Bayesian Data Analysis*. Chapman and Hall/CRC, 3rd edition, 2013. ISBN 978-1439840955.

S. Gillijns, O.B. Mendoza, J. Chandrasekar, B.L.R. De Moor, D.S. Bernstein, and A. Ridley. What is the ensemble kalman filter and how well does it work? In *American Control Conference, 2006*, page 6, June 2006. doi: 10.1109/ACC.2006.1657419.

A. Girard, C. E. Rasmussen, J. Quiñonero Candela, and R. Murray-Smith. Gaussian Process Priors With Uncertain Inputs Application to Multiple-Step Ahead Time Series Forecasting. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems 15*, pages 545–552. MIT Press, 2003.

N. J. Gordon, D. J. Salmond, and A. F. M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Proceedings F (Radar and Signal Processing)*, 140(2):107–113, 1993.

M. S. Grewal, L. R. Weill, and A. P. Andrews. *Global Positioning Systems, Inertial Navigation, and Integration*. Wiley, 2007. ISBN 978-0-470-09971-1.

T. Heine, M. Kawohl, and R. King. Robust Model Predictive Control using the Unscented Transformation. In *2006 IEEE Conference on Computer Aided Control System Design*, pages 224–230, Oct 2006. doi: 10.1109/CACSD-CCA-ISIC.2006.4776650.

R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 1990. ISBN 978-0-521-38632-6.

K. Ito and K. Xiong. Gaussian Filters for Nonlinear Filtering Problems. *IEEE Transactions on Automatic Control*, 45(5):910–927, 2000. doi: 10.1109/9.855552.

T. Jiang, N.D. Sidiropoulos, and G.B. Giannakis. Kalman filtering for power estimation in mobile communications. *Wireless Communications, IEEE Transactions on*, 2(1):151–161, 2003. doi: 10.1109/TWC.2002.806386.

S. J. Julier, J. K. Uhlmann, and H. F. Durrant-Whyte. A New Method for the Nonlinear Transformation of Means and Covariances in Filters and Estimators. *IEEE Transactions on Automatic Control*, 45(3):477–482, 2000.

G. Kitagawa. Monte Carlo Filter and Smoother for Non-Gaussian Nonlinear State Space Models. *Journal of Computational and Graphical Statistics*, 5(1):1–25, 1996.

J. Ko, D. J. Klein, D. Fox, and D. Haehnel. GP-UKF: Unscented Kalman Filters with Gaussian Process Prediction and Observation Models. In *Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on*, volume 54, pages 1901–1907. IEEE, 2007. doi: 10.1109/IROS.2007.4399284.

X. R. Li and Z. Zhao. Measuring Estimator's Credibility: Noncredibility Index. In *Information Fusion, 2006 9th International Conference on*, pages 1–8, 2006. doi: 10.1109/ICIF.2006.301770.

J. McNamee and F. Stenger. Construction of fully symmetric numerical integration formulas. *Numerische Mathematik*, 10(4):327–344, 1967. ISSN 0029599X. doi: 10.1007/BF02162032.

T. P. Minka. Deriving Quadrature Rules from Gaussian Processes. Technical report, Statistics Department, Carnegie Mellon University, Tech. Rep, 2000.

C. J. Oates, M. A Osborne, and M. Girolami. Frank-Wolfe Bayesian Quadrature : Probabilistic Integration with Theoretical Guarantees. *ArXiv e-prints*, pages 1–19, 2015. URL `http://arxiv.org/abs/1506.02681`.

A. O'Hagan. Bayes-Hermite quadrature. *Journal of Statistical Planning and Inference*, 29(3):245–260, 1991. doi: 10.1016/0378-3758(91)90002-V.

M. A. Osborne and R. Garnett. Bayesian Quadrature for Ratios. In *International Conference on Artificial Intelligence and Statistics*, pages 832–840, 2012.

M. A. Osborne, C. E. Rasmussen, D. K. Duvenaud, R. Garnett, and S. J. Roberts. Active Learning of Model Evidence Using Bayesian Quadrature. In *Advances in Neural Information Processing Systems (NIPS)*, pages 46–54, 2012.

H. Poincaré. *Calcul des probabilités*. Paris, Gauthier-Villars, 1896.

J. Prüher and M. Šimandl. Bayesian Quadrature Variance in Sigma-point Filtering. In J. Filipe, O. Gusikhin, Madani K., and J. Sasiadek, editors, *Informatics in Control, Automation and Robotics, 12th International Conference, ICINCO 2015 Colmar, Alsace, France, 21-23 July, 2015 Revised Selected Papers*, volume 370 of *Lecture Notes in Electrical Engineering*. Springer International Publishing, 2016.

C. E. Rasmussen and Z. Ghahramani. Bayesian Monte Carlo. In S. Becker, S. Thrun, and K. Obermayer, editors, *Advances in Neural Information Processing Systems 15*, number 1, pages 505–512. MIT Press, 2003.

C. E. Rasmussen and C. K. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006. ISBN 978-0-262-18253-9.

I. M. Ross, R. J. Proulx, and M. Karpenko. Unscented guidance. In *2015 American Control Conference (ACC)*, pages 5605–5610, July 2015. doi: 10.1109/ACC.2015.7172217.

F. Sandblom and L. Svensson. Moment Estimation Using a Marginalized Transform. *IEEE Transactions on Signal Processing*, 60(12):6138–6150, 2012.

S. Särkkä. *Bayesian Filtering and Smoothing*. Cambridge University Press, New York, 2013. ISBN 978-1-107-61928-9.

S. Särkkä, J. Hartikainen, L. Svensson, and F. Sandblom. Gaussian Process Quadratures in Nonlinear Sigma-Point Filtering and Smoothing. In *Information Fusion (FUSION), 2014 17th International Conference on*, pages 1–8, 2014.

S. Särkkä, J. Hartikainen, L. Svensson, and F. Sandblom. On the relation between Gaussian process quadratures and sigma-point methods. *Journal of Advances in Information Fusion*, 11:31–46, June 2016. ISSN 1557-6418. URL `http://arxiv.org/abs/1504.05994`.

G. L. Smith, S. F. Schmidt, and L. A. McGee. Application of statistical filter theory to the optimal estimation of position and velocity on board a circumlunar vehicle. Technical report, NASA Ames Research Center: Technical Report R-135, 1962.

F. Tronarp, R. Hostettler, and S. Särkkä. Sigma-point filtering for nonlinear systems with non-additive heavy-tailed noise. In *2016 19th International Conference on Information Fusion (FUSION)*, pages 1859–1866, July 2016.

L. Wasserman. *All of Nonparametric Statistics*. Springer, 2007. ISBN 978-0387251455.

Y. Wu, D. Hu, M. Wu, and X. Hu. A Numerical-Integration Perspective on Gaussian Filters. *IEEE Transactions on Signal Processing*, 54(8):2910–2921, 2006.

H. Zangl and G. Steiner. Optimal Design of Multiparameter Multisensor Systems. *IEEE Transactions on Instrumentation and Measurement*, 57(7):1484–1491, July 2008. ISSN 0018-9456. doi: 10.1109/TIM.2008.917175.